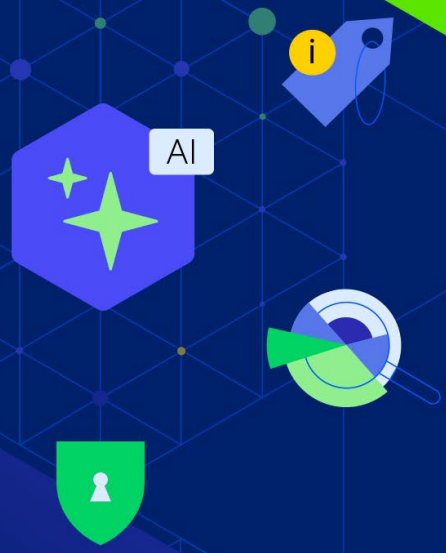
Progress® Data Platform

If Your AI Isn't Boring, It Isn't Ready: Reliable, Governed AI



WHITEPAPER



Why Boredom Matters More Than You Think

Boredom has a bad reputation, especially in the enterprise and AI space. But it's one of the most useful signals humans have. It shows up when intense focus isn't needed, usually when what we're doing is repetitive, low-value or no longer meaningful. This drives us to change course. Researchers often describe boredom as a self-regulatory emotion, a nudge to redirect our efforts towards something more rewarding.

It also has a strange upside. When external stimulation drops our minds don't turn off, they turn inward. Mind-wandering and internal recombination can increase, and in the right conditions that can produce new ideas and better problem-solving. Even experimental work has shown that being bored increases creativity since boredom creates space for daydreaming and innovation.

So boredom isn't a failure mode. It's evolution's way of saying: stop chasing noise and start seeking value. And this is exactly why "Make AI Boring" matters.

Because the future of enterprise AI isn't just reactive chat. It's proactive systems, agents that act on behalf of users. But here's the catch—proactive AI only becomes valuable when it becomes boring first, when it's reliable, verifiable and no longer requires intense oversight But if an agent is unleashed before it's trusted, it doesn't feel like magic. It feels risky. Because in the enterprise, the path to autonomy runs directly through reliability.



© 2026 Progress. All Rights Reserved.

Executive Summary

AI feels like transformative technology that comes along once in a generation. It writes, summarizes, reasons and automates work at a speed that can make entire business processes feel obsolete overnight. That's why the phrase "Make AI Boring" can sound wrong, almost belittling—like if AI is revolutionary, why would we want it to be boring?

Because in business, the most valuable systems don't feel revolutionary day to day. They feel dependable. The moment AI becomes truly useful in the enterprise is when it bridges into the territory of being governable, measurable, secure and operationally consistent. This is "boring" in the best sense: It works the same way every time, for the same reasons, under the same controls and it can be trusted when outcomes matter.

This whitepaper explains why "boring" is the missing prerequisite for scaling AI beyond pilots, why the market is shifting from reactive AI to proactive agent frameworks and how organizations can earn the right to safe autonomy through reliability. It also outlines Progress' approach of leveraging the Progress Data Platform, semantic enrichment and governed retrieval patterns to turn AI from exciting demos into dependable business outcomes.

The path to enterprise AI value is not more intelligence. It's reliability that's earned over time, proven in workflows and enforced through governance.



© 2026 Progress. All Rights Reserved.

The Shock: Why "Boring AI" Sounds Wrong

When people hear boring, they think uninspired or behind the times. AI, by contrast, is marketed as the opposite: dramatic breakthroughs, jaw-dropping demos, increased capability curves. So, calling AI boring can trigger an emotional recoil: Why would we diminish the most exciting tool we've ever had? The misunderstanding is that boring refers to capability but instead refers to **behavior**.

Enterprise buyers don't measure technology by how impressive it is in a demo. They measure it by whether it behaves consistently when it touches real operations, customer experience, compliance, safety, financial exposure, brand risk and daily productivity. In other words, adoption isn't won by excitement; it's won by **confidence**. And confidence is built through repetition, continuity and the slow accumulation of proof.

That's why "flashy AI" can backfire. A team ships something impressive, usage spikes, a handful of unpredictable errors appear, and the initiative gets labeled "too risky." The business doesn't say: "AI failed." It says something worse: "AI can't be trusted here." And that sentiment is hard to reverse.

Trust Debt: The Hidden Cost of Unreliable AI

A useful way to think about early-stage AI is that it creates a new kind of debt: **trust debt**. In fast innovation cycles, instability is often tolerated because capability is improving so quickly. But enterprise environments behave differently. They reward continuity and punish surprises. And they treat "mostly right" as "not good enough" in high-stakes contexts.

Trust debt shows up in subtle, but costly ways: the same question can yield meaningfully different answers. The model can sound confident when it's wrong. Outputs can't be traced back to source evidence, or updates change behavior without any guardrails. The result is predictable: teams spend more time validating outputs than acting on them.

When this happens, AI doesn't just fail to deliver value, it creates a new workload that requires verification, correction, escalation, audit defense and stakeholder reassurance.

© 2026 Progress. All Rights Reserved.

That's why many pilots stall. Not because the model is weak, but because the system isn't dependable enough to earn trust.

This is where "Make AI Boring" becomes practical. Boring here means the system behaves as expected, under control and with backing evidence. It becomes something teams can rely on like they do other enterprise systems. Nobody wants to be caught off guard by a database, identity layer or billing system.

Feature	"Exciting" AI (The Lab)	"Boring" AI (The Enterprise)
Output	Creative, unpredictable, poetic	Consistent, grounded, factual
Process	Black box reasoning	Transparent, evidence-based retrieval
Control	Autonomous and ungoverned	Policy-driven and governed
Failure Mode	Unverifiable hallucinations and risk	Graceful degradation
Business Value	Viral social posts	Operational ROI and scale

From Reactive AI to Proactive Agents

Most organizations are currently using AI in a **reactive** posture: a user asks, the model answers. Even when retrieval is involved, the workflow is still user-led. The human decides what to ask, what to believe and what to do next.

But the next phase is increasingly **proactive**. Systems will anticipate needs, agents will orchestrate steps across tools and AI will increasingly recommend, and eventually execute, actions inside defined boundaries. This shift changes everything, because when AI moves from "answering" to "acting," the cost of error rises dramatically. A wrong answer can be ignored; a wrong action can cause operational and reputational damage.

That's why proactive agents introduce a hard requirement: **deep sustained trust**. If users feel they must double-check the agent every time, autonomy becomes a liability. It creates anxiety instead of value. Proactivity only becomes an advantage when users believe the agent is reliably correct in its domain, and when it's uncertain, it behaves safely.

© 2026 Progress. All Rights Reserved.



As boring as this might sound, sometimes the best answer an AI can give is no.

This is the heart of "proactively reliable" AI: when reliability is not an add-on. It is the prerequisite for safe autonomy.

To make "boring AI" operational rather than rhetorical, define it as properties a system can actually deliver and measure. These are:

It's predictable.
It behaves consistently given similar inputs and contexts. It doesn't swing wildly with phrasing. It degrades gracefully under uncertainty.

It's evidence-based.
It can show what sources were used, why they were selected and how conclusions were formed, especially for regulated or high-impact workflows.

It's governed.
It respects policy, access controls and workflows for approval and escalation. It's clear what it is allowed to do and certain on what it must not do.

It's operational.
It's observable, maintainable and controllable like other enterprise systems by providing:

- Monitoring and alerting
- Audit logs
- Measurable SLAs/SLOs
- Change management for model/prompt updates
- Rollback paths

You can still build innovative experiences on top. But the core behavior, the part that touches enterprise risk, should feel stable, not surprising.

© 2026 Progress. All Rights Reserved.

Use Case Highlight:

Modernization That Prioritizes Compliance, Continuity and Future AI Readiness (Mississippi Division of Medicaid Case Study)

A public-sector healthcare organization faced a modernization trap: legacy systems couldn't keep pace with interoperability mandates, procurement timelines were slow and compliance requirements were strict. And a rip-and-replace configuration carried too much risk and too much cost.

They adopted an incremental modernization strategy using the Progress Data Platform, prioritizing standards compliance, transparency and reusable architecture over novelty. This reduced modernization cost materially versus an original large-scale approach. It also eliminated the change-order spiral that often derails complex public programs. At the same time, the organization positioned itself for future AI use cases because the foundation was now standards-based, governable and adaptable.

This is "Make AI Boring" in the most important sense: doing the unglamorous work that facilitates systems to remain trustworthy even as requirements change.

What made it boring, yet valuable:

- Standards-based interoperability and audit-ready transparency
- Incremental rollout protecting continuity of service
- Reusable architecture enabling future use cases safely
- Built-in (not bolted-on) governance and compliance

In mission-critical environments, "move fast" means build foundations that don't break—then innovate safely on top.

© 2026 Progress. All Rights Reserved.

A Maturity Model for Proactively Reliable AI

The mistake many teams make is trying to jump to "agentic" AI before they've earned trust. The smarter move is to progress through stages, expanding autonomy only when reliability is proven.

It starts with reliable answers, grounded in enterprise sources, supported by citations and provenance and tested against evaluation sets so quality doesn't drift unnoticed. From there, teams introduce assisted actions: the AI suggests, drafts and prepares work, but humans approve before anything changes in a system of record. Only after a track record is established does autonomy expand into low-stakes actions where mistakes can be overcome and the blast radius is small. Over time, and with strong identity, policy controls, monitoring and change management, autonomy can broaden safely.

Don't ask customers to trust autonomy. Prove reliability first, then expand scope.

This approach is not slower. It's an intentional way to prevent the trust collapse that kills adoption. The fastest way to kill agentic adoption is to start with high-stakes autonomy before the system has earned trust. The practical path is staged, expanding autonomy only when reliability is proven.

Stage 1: Foundational, reliable answers

At this stage, the goal is defensibility—not delight. Teams focus on:

- Grounding answers in authoritative enterprise sources
- Improving retrieval quality with context and semantics
- Producing evidence packages (answer + sources + rationale)
- Measuring quality against reference sets to drive consistency

Stage 2: Assisted, suggested actions

The AI proposes actions, but humans approve. This reduces risk while building confidence by:

- Drafting a response, but requiring review
- Proposing workflow steps without triggering them
- Generating summaries without publishing them
- Logging the rationale for what it would have done and why

Stage 3: Low-stakes, bounded and reversible autonomy

Autonomy begins where mistakes are recoverable and the blast radius is small when:

- Routine tasks have built-in rollback
- Constraints are placed on actions with explicit policy checks
- Clear auditing and escalation pathway is available to humans

Stage 4: Expanded, policy-driven autonomy

Only after reliability has been sustained can:

- Progressive permissions be set up by role, context and risk
- Continuous evaluation and drift monitoring happen
- Strong controls be triggered around updates and behavioral changes

© 2026 Progress. All Rights Reserved.

Use Case Highlight:

Legal Research That Gets Faster and More Defensible

A global, regulated professional services organization faced a familiar challenge: legal associates were spending hours compiling legislation, case law, internal guidance and prior analyses. The work was slow, costly and risk-prone, not only because information was fragmented, but because the hardest part was establishing what sources were authoritative and making sure answers could be defended.

They built a centralized, governed knowledge layer using the Progress Data Platform, then layered an AI legal assistant on top. The design goal wasn't to be impressive; it was to be reliably correct through:

- Continuous content enrichment with semantic/contextual signals to reduce ambiguity
- Hybrid retrieval approach, with vector fallback when applicable, to gather the best evidence
- Sustained trust with the systematic measuring of AI outputs against human-created references

The result was a dramatic time savings, as hours of legal associate work was condensed to minutes, without compromising confidence. The organization didn't just accelerate research, it created a repeatable approach for scaling into other regulated, knowledge-intensive workflows.

What made it boring, yet valuable:

- Evidence-first answers grounded in authoritative sources
- Semantic context to improve retrieval precision and consistency
- Continuous validation loops to maintain quality over time
- Consistent output governance that went from plausible to defensible

In regulated work, "AI that sounds right" is a liability; AI that can prove "why it's right" becomes infrastructure.

© 2026 Progress. All Rights Reserved.

Progress' Approach to Governed, Grounded Enterprise AI

A common failure mode in enterprise AI is assuming the model is the product. In reality, the model is only one component. Reliability is usually determined by the surrounding system: context, retrieval, semantics, governance and operations.

The Progress Data Platform approaches "Make AI Boring" as an enterprise outcome: trusted AI systems grounded in enterprise truth, enriched with semantics and delivered with governance and operational control.

At a high level, the Progress Data Platform approach centers on three ideas:

1

Build an AI-ready knowledge foundation

Instead of treating enterprise content like generic internet text, the platform focuses on turning enterprise data into something AI can use reliably to:

- Unify and access data across systems
- Connect content, records and context
- Preserve provenance and lineage
- Maintain consistent governance and access controls

This is where a platform approach matters. If AI is stitched together as disposable layers, trust is fragile. When AI is grounded in a durable data platform, trust compounds.

2

Use semantics to improve grounding and explainability

"Boring AI" requires more than vector similarity. It requires meaning and structure to:

- Enrich content with semantic signals
- Use taxonomies/ontologies to standardize interpretation
- Apply knowledge-graph relationships to improve precision and explainability
- Support hybrid retrieval patterns that balance recall, accuracy and control

This is how outputs become defensible, because the system can explain not only what it found, but why it's relevant and how it connects to enterprise truth.

© 2026 Progress. All Rights Reserved.

3

Deliver a governed path toward agentic capabilities

The Progress Data Platform view of agentic AI is pragmatic: you expand autonomy carefully, starting with low stakes and high reversibility, and only move into action when reliability is proven and trust is earned.

In practice, that means enabling agentic experiences that:

- Begin with suggestions
- Include approvals and policy checks
- Log actions and evidence
- Provide rollback paths
- Evolve toward automation

Our message here is consistent: the goal is not having agents everywhere; the goal is having proactively reliable agents that operate safely inside enterprise constraints.



© 2026 Progress. All Rights Reserved.

Use Case Highlight:

News Agency Makes Discovery Predictable Across Languages & Formats

A national news agency distributes multilingual content across text, photos, video and graphics under constant time constraints. Discovery depended on manual tagging, often done by authors under deadline. This created inconsistencies that degraded search relevance and customer experience.

Instead of bolting generative AI (GenAI) onto unreliable content foundations, the agency built a semantic knowledge layer using the Progress Data Platform. Meaning became explicit through ontologies, taxonomies and automated tagging across structured and unstructured content. Enrichment integrated directly into CMS and DAM workflows, so metadata quality wasn't optional, it was operationalized.

The outcome was a calmer, more reliable discovery experience where:

- Customers found relevant stories faster and more consistently
- Internal teams spent less time re-tagging and chasing missing context
- The agency gained a stable foundation for new digital product experiences

What made it boring, yet valuable:

- Consistent multilingual semantics embedded into the workflow
- Automated enrichment replacing inconsistent manual tagging
- Relationship-driven metadata improving relevance predictably
- Stable platform foundation to support new products

For many organizations, the most valuable AI capability isn't generation.
It's reliable discovery under real-world constraints.

© 2026 Progress. All Rights Reserved.

Implementation Checklist for Enterprise AI Governance

If you want to operationalize "Make AI Boring," treat it as a program with explicit deliverables like these:

Define the trust boundary

Establish where AI is allowed to operate and where it isn't, including:

- Workflows that are in scope
- Sources that are authoritative
- Requirements for human approval
- Red lines for regulated decisions and sensitive actions

Establish quality measurement

Treat AI quality like software quality with:

- Evaluation sets per use case
- Regression tests across model/prompt changes
- Monitoring for drift and failure modes
- Thresholds for reliability before expanding autonomy

Instrument evidence and traceability

Make evidence a first-class requirement by including:

- Source references and provenance
- Retrieval rationale (i.e., why these sources?)
- Audit logs for outputs and actions
- Clear accountability paths

Expand autonomy carefully

Start with reversibility as follows:

- Suggestion > approval > low-stakes automation > expansion
- Action constraint with policy and least privilege
- Real and tested rollback and escalation paths

Key insight: The fastest way to scale AI is to scale trust mechanisms, not experiments.

© 2026 Progress. All Rights Reserved.

The Most Transformative AI Will Feel the Least Dramatic

The future of enterprise AI is not an endless parade of impressive demos. It's the quiet emergence of AI as dependable infrastructure embedded into workflows, governed by policy, grounded in enterprise truth and trusted enough to act.

That's what **Make AI Boring** really means: to stop optimizing for novelty, start optimizing for reliability, earn the right to autonomy and turn AI into something the business can depend on without fear.

Progress' approach is aligned to this reality: to build a durable, semantic, governed data foundation with the Progress Data Platform. Then employ that foundation to deliver reliable AI outcomes today and a safe, proven path toward proactive agentic capabilities tomorrow.

Turn AI from Exciting Demos into Dependable Infrastructure

Learn why "boring AI" is a prerequisite for safe autonomy. Then see how the Progress Data Platform supports governed retrieval, semantic context and enterprise control.

© 2026 Progress. All Rights Reserved.

FAQ

What does “boring AI” mean in the enterprise?
Boring AI is AI you can rely on. It behaves predictably, is grounded in authoritative enterprise sources, produces defensible outputs with clear evidence and operates within governance controls (e.g., access, policy, audit logs, approvals). It’s not less capable; it’s more disciplined.

Why is “boring” the prerequisite for agentic or proactive AI?
Because proactive systems don’t just answer questions, they can recommend and execute actions. The cost of error rises sharply when AI touches systems of record, compliance processes or customer-facing workflows. If users feel they must verify every action, autonomy creates anxiety rather than value. Reliability earns the right to be proactive.

What’s the difference between reactive AI and proactive agent frameworks?
Reactive AI responds when a user asks (ask > answer). Proactive agent frameworks sense context, recommend next steps, orchestrate tasks across tools and can take actions (sense > decide > act) within defined boundaries. The proactive shift demands stronger governance, traceability and control.

How does grounded AI reduce hallucinations?
Grounded AI ties responses to trusted enterprise sources rather than relying on probabilistic guesswork. In practice, it uses retrieval strategies, like retrieval-augmented generation (RAG), to pull relevant evidence, then generates answers that reference those sources. The result is fewer hallucinations, clearer accountability and defensible outputs.

What is RAG, and why isn’t “vector-only” retrieval enough?
Retrieval-augmented generation (RAG) combines search/retrieval with generation, so answers are based on evidence. Vector-only retrieval can be useful, but it can also return similar content that’s not authoritative, current or relevant to policy constraints. Enterprise-grade RAG typically needs hybrid retrieval (e.g., keyword + semantic / vector + sometimes graph/metadata filters) to improve precision, control and repeatability.

What is semantic RAG, and how does it improve reliability?
Semantic RAG adds meaning and structure, taxonomies, ontologies, entity resolution and relationship models (including knowledge graphs, where valuable), to improve retrieval accuracy and consistency. It reduces ambiguity (e.g., what “customer” means), improves explainability and helps keep outputs stable as data grows.

How do you make AI outputs defensible for regulated or high-impact workflows?
Defensible outputs require an evidence trail. That usually means source citations, provenance (where the data came from), timestamps, retrieval rationale and audit logs, plus evaluation against reference sets to measure accuracy over time. For higher-risk actions, you can add approvals and human-in-the-loop checks.

What governance controls do proactive AI agents require?
At a minimum: identity and access controls, policy enforcement (what the agent can/can’t do), audit logs for decisions and actions, approval workflows for higher-risk steps and rollback paths for recoverability. As autonomy expands, continuous monitoring for drift and incident response becomes essential.

© 2026 Progress. All Rights Reserved.

How do you "earn the right" to autonomy without increasing risk?

Start with reliability, then expand in stages: (1) grounded, evidence-based answers; (2) assisted actions where the AI makes suggestions and humans approve; (3) low-stakes autonomy with small blast radius and rollbacks; and (4) broader autonomy governed by policy, monitoring and change control. Then expand scope only after sustained performance.

How do we measure whether our AI is becoming reliable, or "boring"?

Look for workflow signals: less time spent verifying outputs, higher repeatability for the same intent, fewer escalations, stronger audit readiness with standard evidence, sustained adoption after novelty fades and gradual expansion of autonomous scope. This is due to the expansion of trust, not the weakening of controls.

What are the most common reasons enterprise AI pilots stall?

Usually trust debt: inconsistent outputs, lack of evidence/provenance, poor retrieval quality, unclear governance and operational gaps in terms of monitoring or change control. When AI creates more verification work than value, adoption stalls—even if the model itself is powerful.

What's the first step to making AI boring in our organization?

Define the trust boundary by scoping out use cases, identifying authoritative sources, deciding which actions to permit, determining which requires approval and what shouldn't be automated. Then design evidence and traceability as defaults, not as afterthoughts.

Explore the Progress Data Platform | **Talk to an Expert**

About Progress Software

Progress Software (Nasdaq: PRGS) empowers organizations to achieve transformational success in the face of disruptive change. Our software enables our customers to develop, deploy and manage responsible AI-powered applications and personalized digital experiences with agility and ease. Businesses of all sizes get a trusted provider in Progress, with the products, expertise and vision they need to turn AI disruption into a competitive advantage. Millions of developers and technologists at hundreds of thousands of organizations depend on Progress every day. Learn more at www.progress.com

© 2026 Progress Software Corporation and/or its subsidiaries or affiliates.
All rights reserved. Rev 2026/02 | RITM0349288

Worldwide Headquarters

Progress Software Corporation
15 Wayside Rd, Suite 400, Burlington, MA 01803, USA
Tel: +1-800-477-6473

facebook.com/progresssw
x.com/progresssw
youtube.com/progresssw
linkedin.com/company/progress-software
progress_sw_

Progress